MANAGEMENT SCIENCE
Vol. 38, No. 3, March 1992
Printed in U.S.A.

JOINT COHERENCE IN GAMES OF INCOMPLETE INFORMATION*

ROBERT F. NAU
The Fuqua School of Business, Duke University, Durham, North Carolina 27706

Decisions are often made under conditions of uncertainty about the actions of supposedly-rational competitors. The modeling of optimal behavior under such conditions is the subject of noncooperative game theory, of which a cornerstone is Harsanyi's formulation of games of incomplete information. In an incomplete-information game, uncertainty may surround the attributes as well as the strategic intentions of opposing players. Harsanyi develops the concept of a Bayesian equilibrium, which is a Nash equilibrium of a game in which the players' reciprocal beliefs about each others' attributes are consistent with a common prior distribution, as though they had been jointly drawn at random from populations with commonly-known proportions of types. The relation of such game-theoretic solution concepts to subjective probability theory and nonstrategic decision analysis has been controversial, as reflected in critiques by Kadane and Larkey and responses from Harsanyi, Shubik, and others, which have appeared in this journal. This paper shows that the Bayesian equilibrium concept and common prior assumption can be reconciled with a subjective view of probability by (i) supposing that players elicit each others' probabilities and utilities through the acceptance of gambles, and (ii) invoking a multi-agent extension of de Finetti's axiom of coherence (no arbitrage opportunities, a.k.a. "Dutch books"). However, the Nash property of statistical independence between players is weakened, and the probability distributions characterizing a solution of the game admit novel interpretations.
(ARBITRAGE; COHERENCE; SUBJECTIVE PROBABILITY; NONCOOPERATIVE GAMES; BAYESIAN EQUILIBRIUM; CORRELATED EQUILIBRIUM; COMMUNICATION EQUILIBRIUM; COMMON PRIOR ASSUMPTION)

1. Introduction

In the tradition established by von Neumann and Morgenstern (1944) and extended by Harsanyi (1967), analyses of rational strategic behavior begin with a careful description of the "rules of the game." The rules specify a set of players, a set of strategies available to the players, a set of states of nature representing uncertainty about the *types* (utility functions and information states) of the players, and a common prior probability distribution quantifying that uncertainty. The rules of the game are assumed to be common knowledge among the players, as is the fact that all players are Bayesian rational subjective-expected-utility maximizers, although specific allowances are sometimes made for errors or irrational play. Given this exogenously determined common-knowledge structure, game theory seeks to endogenously determine the strategy or set of strategies that should be played.

Various authors (Armbruster and Böge 1979, Böge and Eisele 1979, Mertens and Zamir 1985, Tan and Werlang 1988) have established that common knowledge of the rules of a game and common knowledge of the players' rationality can be formalized in mathematically tractable ways. By pursuing the mathematics of infinite regress to its limits, it can be shown that a postulated hierarchy of reciprocal beliefs among the players "terminates" at a single order of infinity, and, moreover, the set of such infinite-order beliefs contains belief-closed subsets in which uncertainty about the game's structure is summarized by a common prior distribution over a finite set of types, as proposed by Harsanyi (1967). However, this demonstration that common knowledge is a mathematically well-defined concept does not provide an operational method of determining

* Accepted by Irving H. LaValle and the Editors of the Special Issue on Group Decision and Negotiation; received February 5, 1990. This paper has been with the author 3 months for 3 revisions.

0025-1909/92/3803/0374$01.25

whether or to what extent it is achieved in a specific game situation: the assumptions on which it rests are beyond possibility of empirical verification.

The objects of common-knowledge assumptions are subjective attributes of the players, namely their personal probability distributions and utility functions. It is well known that such attributes are difficult to elicit in practice, even from isolated individuals in nonstrategic situations, and decision theorists have increasingly questioned whether the underlying assumptions of subjective expected utility theory are appropriate even as idealizations of rational behavior (e.g., Aumann 1962, Machina 1982, Fishburn 1982, Seidenfeld 1988, Schmeidler 1989). Problems of elicitation are magnified where strategic interactions are present or where probability and utility are measured jointly: for example, strategic considerations may affect the precision with which probabilities can be measured (Ellsberg 1961, Leamer 1986, Nau 1992), and unobservability of prior wealth distributions may render it difficult to separate personal probabilities from utilities (Kadane and Winkler 1988). For these reasons it is of interest to determine the extent to which game-theoretic concepts of common knowledge and mutually expected rationality can be reformulated in terms of first-order observable behavior without presupposing the existence of sharply-defined hierarchies of reciprocal beliefs.

This paper describes an operational method through which, in principle, the players might achieve common knowledge of the subjective parameters of a noncooperative game, and which leads to an extremely simple characterization of mutually expected rationality. It is assumed that there is a *set of events* representing possible outcomes of a game and a *market* in which monetary claims (lottery tickets or gambling contracts) can be written and enforced on those events. The market serves as the medium of communication among the players and is interpreted as a canonical model of the institutional or cultural context of the game. The parameters of the game—including the players' utilities, probabilities, and the structures of their strategy sets and information partitions—are subjectively revealed through transactions in the market prior to the official start of play. Thus, preplay communication is to some extent endogenized. The market framework allows a natural *primal* characterization of rational preplay communication and intraplay decision-making to be given in terms of observable behavior. Game-theoretic concepts, such as common prior distributions, incentive-compatibility constraints, and equilibria among strategy selections, will be shown to emerge as elements of a *dual* characterization.

The behavior of the players is defined to be strategically rational if the outcome of the game, in the context of the market which defines its rules, creates no arbitrage opportunities. This can be viewed as a multi-agent extension of the axiom of *coherence* proposed by de Finetti (1937, 1974) as a basis for subjective probability theory; hence it is dubbed *joint coherence*. The concept of joint coherence will be applied here to games of incomplete information, extending results previously obtained for games of complete information by Nau and McCardle (1990). In the complete-information case, the dual characterization of joint coherence was shown to be a correlated generalization of Nash equilibrium, namely Aumann's (1974, 1987) concept of *correlated equilibrium*. In the case of incomplete-information games without observable communication (i.e., without mechanical coordination devices), the dual characterization will be shown to be a correlated generalization of Harsanyi's (1967) *Bayesian equilibrium* concept. In games with observable communication, the dual characterization will be shown to be the concept of *communication equilibrium* (Myerson 1985, Forges 1986).

These results provide a constructive resolution of the controversy between "subjective" and "game-theoretic" views of rational decision-making under uncertainty that was sparked in this journal by the papers of Kadane and Larkey (1982,[1] 1983) and the

[1] "[W]e do not understand the search for solution concepts that do not depend on the beliefs of each player about the others' likely actions and yet are so compelling that they will become the obvious standard of play for all those who encounter them."

responses of Harsanyi (1982), Shubik (1983), Kahan (1983), Roth and Schoumaker (1983), and Rothkopf (1983). It is seen that there is no essential difference between strategic and nonstrategic rationality: both are characterized by the same criterion of coherence, appropriately tailored to agents' endowments of information and control. This criterion, otherwise known as the "arbitrage principle," also characterizes rational competitive behavior in a variety of other economic contexts (Nau and McCardle 1991).

Throughout this paper, it will be assumed for the sake of brevity and simplicity that the players' marginal utilities for money are constant across outcomes of the game, as though they were risk-neutral, so that their true probabilities and comparative utilities for outcomes can be directly elicited via the acceptance of monetary gambles. However, the results obtained here generalize in an interesting way to the more realistic case of risk aversion and state-dependent marginal utilities. The nature of this generalization will be sketched briefly in the concluding section and is treated in more detail by Nau (1990, 1991a, b).

2. The Rules of the Game

Consider a finite noncooperative game among N players, and let Ω denote its set of distinct outcomes. Ω will be assumed to be objectively given (exogenously determined) in the sense that it provides a basis for enforceable gambling contracts. All other aspects of the game—including its payoffs, information partitions, and probability distributions on states or strategies—will be considered to be subjectively revealed (endogenously determined) by the players. To conform with standard game-theoretic notation, the outcome set will be decomposed as $\Omega = S \times T$, where S is interpreted as a set of joint strategies available to the players and T is interpreted as a set of states of nature representing exogenous uncertainty.[2] The strategy set in turn will be decomposed as $S = S_1 \times \cdots \times S_N$, where S_n is interpreted as the strategy set of player n. Similarly, the state set will be decomposed as $T = T_1 \times \cdots \times T_N$, where the elements of T_n are interpreted as information states (types) of player n, upon which his choice of strategies may be conditioned. $s = (s_1, \ldots, s_N)$ and $t = (t_1, \ldots, t_N)$ will denote elements of S and T, respectively. Finally, S_{-n} and T_{-n} will denote the sets of strategies and types available to all players other than n, with generic elements s_{-n} and t_{-n}, respectively.

The game is "noncooperative" insofar as the players may be unable or unwilling to communicate or to enter binding contracts concerning the strategies they will play. However, the environment is not devoid of communication or contractual obligations: it is assumed that money is available as a medium of exchange; that there is contingent-claims market in which monetary gambling contracts can be written and enforced with respect to outcomes of the game; and that this is the mechanism through which the rules of the game become common knowledge. The market assumption merely extends de Finetti's (1937) operational definition of probability to the situation in which two or more agents elicit each others' probabilities simultaneously through gambling.

DEFINITIONS. A *gamble* is a monetary payoff function (vector) defined on the outcome set $S \times T$. A gamble $\mathbf{g}$ is *acceptable* to some player if he asserts that he is willing to receive the payoff $\beta g(s, t)$ when strategy $s \in S$ is played and state $t \in T$ obtains, where β is any small[3] nonnegative number chosen by an opponent (*e.g.*, an observer of the game or another player) after the announcement of $\mathbf{g}$ but prior to the realization of s

[2] This decomposition of the outcome set is itself subjective. The extent to which outcomes of the game are believed to be controlled by different players or by nature will ultimately be revealed through the structure of gambles accepted by the players.

[3] Note that acceptance of a gamble commits a player to transactions whose payoffs are only small multiples of those of the gamble itself, hence the scale of an acceptable gamble is arbitrary. If necessary for technical reasons—e.g., to explicitly accommodate strict risk aversion—"small" can be defined to mean "infinitesimal" (Nau 1991a).

and t. $\mathbf{g}_1$ *dominates* $\mathbf{g}_2$ if $g_1(s, t) \geq g_2(s, t)$ for all (s, t). The set of all gambles acceptable to one or more players will be denoted $\mathcal{A}$. $\mathbf{0}$ and $\mathbf{1}$ will denote payoff vectors whose elements are identically 0 and 1, respectively. The symbols $\mathbf{E}$ and $\mathbf{F}$ will denote arbitrary events (subsets of $S \times T$) and also the indicator functions thereof. That is, $E(s, t) = 1$ if $\mathbf{E}$ is true in outcome (s, t), and $E(s, t) = 0$ otherwise. The symbol $\mathbf{1}_{njr}$ will denote the particular event that player n is of type r and plays strategy j. That is, $1_{njr}(s, t) = 1$ if $s_n = j$ and $t_n = r$, and $1_{njr}(s, t) = 0$ otherwise. Addition and multiplication are defined pointwise. Thus, $\alpha_1\mathbf{g}_1 + \alpha_2\mathbf{g}_2$ is the payoff function whose value is $\alpha_1 g_1(s, t) + \alpha_2 g_2(s, t)$ in outcome (s, t), and $(\mathbf{E} - p)\mathbf{F}$ is the function whose value is $(E(s, t) - p)F(s, t)$ in outcome (s, t).

The players' beliefs and preferences are revealed to each other by the gambles they accept in the market. For example, if p is a number between 0 and 1 and $\mathbf{E}$ is an event, then acceptance of the gamble $\mathbf{E} - p$ indicates a belief that the probability of $\mathbf{E}$ is at least p. If $\mathbf{x}$ and $\mathbf{y}$ are lotteries (arbitrary payoff functions on $S \times T$) and $\mathbf{F}$ is an event, then acceptance of the gamble $(\mathbf{x} - \mathbf{y})\mathbf{F}$ indicates that $\mathbf{x}$ is preferred to $\mathbf{y}$ given knowledge that $\mathbf{F}$ has occurred.

The following structural assumptions are imposed on the set $\mathcal{A}$ of acceptable gambles:

A1 (Dominance). $\mathbf{g}$ dominates $0 \Rightarrow \mathbf{g} \in \mathcal{A}$;

A2 (Linearity). $\mathbf{g} \in \mathcal{A} \Rightarrow \alpha\mathbf{g} \in \mathcal{A}\ \forall\alpha > 0$;

A3 (Additivity). $\mathbf{g}_1, \mathbf{g}_2 \in \mathcal{A} \Rightarrow \mathbf{g}_1 + \mathbf{g}_2 \in \mathcal{A}$.

In other words, gambles are infinitely divisible, additive, and measured in a common currency of which more is preferred to less. The set $\mathcal{A}$ is therefore a convex cone which includes the nonnegative orthant. If transactions based on acceptable gambles are understood to be small, these assumptions are consistent with the axioms of subjective expected utility, but weaker: they do not require subjective probability distributions and utility functions to be separable nor uniquely determined (Nau 1991a).

The subjective structure of the game is assumed to be revealed through two kinds of gambles accepted by the players, *preference gambles* and *belief gambles*, which generate the set $\mathcal{A}$ via A1–A3. Preference gambles reveal the relative differences in payoffs the players perceive between the strategies they choose and those they do not choose, given their information. This idea is formalized as follows: for any $s \in S$ let $s_{-n}j = (s_1, \ldots, s_{n-1}, j, s_{n+1}, \ldots, s_N)$. That is, $s_{-n}j$ denotes the situation in which players other than n adhere to the joint strategy s while player n chooses his jth strategy. Now let $\mathbf{u}_n : S \times T \mapsto \Re$ denote a hypothetical payoff function for player n, expressed in units of personal utility. Following Nau and McCardle (1990), let $\mathbf{u}_{nj} : S \times T \mapsto \Re$ denote the corresponding function specifying the payoffs the player *would have received* by playing strategy j, as a function of nature's and his opponents' actions. That is, $\mathbf{u}_{nj}$ is derived from $\mathbf{u}_n$ according to:

$$u_{nj}(s, t) \equiv u_n(s_{-n}j, t) \qquad \forall(s, t) \in \Omega. \tag{1}$$

If player n's marginal utility for additional monetary wealth is assumed to be constant across outcomes of the game, then he should be willing to accept the gamble $(\mathbf{u}_{nj} - \mathbf{u}_{nk})\mathbf{1}_{njr}$ for every $j, k \in S_n$ and every $r \in T_n$. This merely affirms that, in the event he is observed to play strategy j given information r, player n prefers the relative payoffs yielded by strategy j (as a function of nature's and his opponents' actions) to those which would have been yielded by any other strategy k.[4] We now assume that for each player there are acceptable gambles which are consistent in this way with *some* underlying payoff

[4] In the event strategy j is chosen by player n, this gamble yields increments of utility proportional to the differences in utility he perceives between strategies j and k. On the assumptions that (i) his chosen strategy maximizes expected utility with respect to some probability distribution over nature's and his opponents' actions, and (ii) he will accept any gamble which yields nonnegative incremental expected utility with respect to the same distribution, it follows that a gamble constructed in this way is acceptable. A decomposed method of eliciting such a gamble is described by Nau (1991a).

function, and that this is the method by which payoff functions are effectively revealed. Formally, we adopt:

A4 (Acceptance of preference gambles). For every n, there exists a function $\mathbf{u}_n : S \times T \mapsto \Re$ such that for every j and $k \in S_n$, and every $r \in T_n$: $(\mathbf{u}_{nj} - \mathbf{u}_{nk})\mathbf{1}_{njr} \in \mathcal{A}$, where $\mathbf{u}_{nj}$ and $\mathbf{u}_{nk}$ are derived from $\mathbf{u}_n$ according to (1).

The preference gambles of player n reveal the structure of his strategy set (S_n) and type set (T_n) as well as his payoff function ($\mathbf{u}_n$). The latter is uniquely determined only up to positive affine transformations and/or the addition of terms which are independent of his own actions. The symbol A will henceforth denote the matrix whose columns, indexed by (n, j, k, r), are the payoff vectors of the acceptable preference gambles in A4.

The subjective description of the game is completed through the acceptance of belief gambles, which directly establish bounds on the players' conditional probabilities for events. For each n, let H_n denote a finite set of index numbers, and for all $h \in H_n$ let $\mathbf{E}_{nh}$ and $\mathbf{F}_{nh}$ denote events; and let p^*_{nh} denote a number between 0 and 1. Then:

A5 (Acceptance of belief gambles). For every n, there exist events $\{\mathbf{E}_{nh}, \mathbf{F}_{nh} \mid \mathbf{h} \in \mathbf{H}_n\}$ and numbers $\{p^*_{nh} \mid h \in H_n\}$ such that: $(\mathbf{E}_{nh} - p^*_{nh})\mathbf{F}_{nh} \in \mathcal{A}$.

Thus, p^*_{nh} is asserted by player n to be a *lower probability*[5] for the conditional event $\mathbf{E}_{nh} \mid \mathbf{F}_{nh}$. Let B henceforth denote the matrix whose columns are payoff vectors of the acceptable belief gambles in A5.

A probability distribution π on $S \times T$ will be called a *supporting* probability distribution for player n if it assigns nonnegative expected value to every belief gamble he accepts—i.e., if it satisfies:

$$\pi^T(\mathbf{E}_{nh} - p^*_{nh})\mathbf{F}_{nh} \geq 0 \qquad \forall h \in H_n,$$

or equivalently:

$$P_\pi(\mathbf{E}_{nh} \mid \mathbf{F}_{nh}) \geq p^*_{nh} \qquad \text{or else} \qquad P_\pi(\mathbf{F}_{nh}) = 0 \qquad \forall h \in H_n,$$

where $P_\pi(\mathbf{E} \mid \mathbf{F})$ denotes the conditional probability (or expectation) of $\mathbf{E}$ given $\mathbf{F}$ under the distribution π.[6] In typical models of incomplete-information games, players are not assumed to reveal probabilities they attach to their own information states or strategies; rather, they are assumed only to reveal their probabilities for information states and/or strategies of their competitors conditional on their own possible information states and/or strategies. That is, the events $\mathbf{E}_{nh}$ and $\mathbf{F}_{nh}$ are typically assumed to be measurable with respect to $S_{-n} \times T_{-n}$ and $S_n \times T_n$, respectively. If this is the case, the belief gambles accepted by any one player will not uniquely determine a supporting probability distribution on $S \times T$ even if his conditional probabilities *are* uniquely determined.

As an example, consider the following game of incomplete information discussed by Myerson (1985): there are two players (1 = row, 2 = column) whose strategy sets are $\{T, B\}$ and $\{L, R\}$, respectively, and two states of nature, $\{a, b\}$, whose "prior" probabilities are 0.6 and 0.4, respectively. The state will be revealed to player 2 before strategies are selected, but not to player 1—i.e., a and b are "types" of player 2. The payoff functions of the players are shown in Table 1. The numbers in each pair of parentheses are the payoffs to 1 and 2, respectively. Let $\Omega = \{aTL, aBL, aTR, \ldots, bBR\}$

[5] Note that the hth gamble in A5 yields a payoff of $1 - p_{nh}$ if both $\mathbf{E}_{nh}$ and $\mathbf{F}_{nh}$ occur, a payoff of $-p_{nh}$ if $\mathbf{F}_{nh}$ occurs but not $\mathbf{E}_{nh}$, and a payoff of 0 if $\mathbf{F}_{nh}$ does not occur. This is merely de Finetti's (1937) operational definition of a subjective conditional probability, as generalized by Smith (1961) to the case of lower and upper probabilities. For our purposes, it will suffice to consider only lower probabilities, since an upper probability for an event can be expressed as a lower probability for the complementary event. The use of lower probabilities rather than "sharp" probabilities allows for indeterminacy in the commonly known beliefs due to strategic interactions in the measurement process—cf. Nau (1992).

[6] That is, for any lottery (arbitrary payoff function) $\mathbf{E}$ and any event (indicator function) $\mathbf{F}$, $P_\pi(\mathbf{E}) \equiv \pi^T\mathbf{E} = \sum_{(s,t)\in\Omega} \pi(s, t)E(s, t)$, and $P_\pi(\mathbf{E} \mid \mathbf{F}) \equiv P_\pi(\mathbf{EF})/P_\pi(\mathbf{F})$ if $P_\pi(\mathbf{F}) > 0$.

TABLE 1

Payoff Functions for Game Example

	a		b	
	L	R	L	R
T	(1, 2)	(0, 1)	(1, 3)	(0, 4)
B	(0, 4)	(1, 3)	(0, 1)	(1, 2)

denote the set of eight distinct outcomes of the game, and let $\mathbf{u}_n$ denote the payoff function of player n—i.e., $u_1(aTL) = 1$, $u_2(aTL) = 2$, etc.

To construct the preference gambles which encode this payoff structure as provided in A4, let $\mathbf{u}_T$ denote the function on Ω whose value in state ω is the payoff player 1 would have received by playing T given that nature and his opponent played in accordance with ω. (Here, for notational simplicity, we omit the subscript for the player's number, writing $\mathbf{u}_T$ instead of $\mathbf{u}_{1T}$.) Thus, for example, $u_T(aTL) = u_T(aBL) = 1$, $u_T(aTR) = u_T(aBR) = 0$, etc. The function $\mathbf{u}_B$ for player 1 and the functions $\mathbf{u}_L$ and $\mathbf{u}_R$ for player 2 are defined similarly. Since player 1 is assumed to have no information about the state of nature when he makes his move, his preference gambles are conditioned only on his own actions. He therefore should accept the preference gamble $(\mathbf{u}_T - \mathbf{u}_B)\mathbf{1}_T$, reflecting his preference for $\mathbf{u}_T$ over $\mathbf{u}_B$ in the event that he is observed to play T. Similarly, he should also accept the preference gamble $(\mathbf{u}_B - \mathbf{u}_T)\mathbf{1}_B$. Player 2 is assumed to know the true state of nature at the time he makes his move, so his preference gambles are conditioned on this information as well as his own actions. He therefore should accept the preference gamble $(\mathbf{u}_L - \mathbf{u}_R)\mathbf{1}_{aL}$, reflecting a preference for $\mathbf{u}_L$ over $\mathbf{u}_R$ when he is observed to play L while knowing that the state is a, and similarly for other strategy-state combinations. The payoff vectors for the preference gambles of both players are assembled in the matrix "A" (Table 2).

The belief gambles for this game must encode the information that the prior probabilities of states a and b are 0.6 and 0.4, respectively, and that player 1 will receive no additional information concerning the state before making his move. In other words, it is asserted (presumably by player 1) that the lower conditional probabilities of a and b are 0.6 and 0.4, respectively, *given any move of player* 1. This is conveyed by the acceptance of four gambles whose payoff vectors are $(\mathbf{1}_a - p_a^*)\mathbf{1}_T$, $(\mathbf{1}_a - p_a^*)\mathbf{1}_B$, $(\mathbf{1}_b - p_b^*)\mathbf{1}_T$ and $(\mathbf{1}_b - p_b^*)\mathbf{1}_B$, where $p_a^* = 0.6$ and $p_b^* = 0.4$. These constitute the matrix "B" (Table 3).

Altogether, the subjective rules of the game (strategy sets, payoff functions, information partitions, and prior probabilities) are encoded in the matrices A and B. The columns of these matrices are the payoff vectors of the acceptable preference and belief gambles defined in A4 and A5, which generate (via A1–A3) the set $\mathcal{A}$ of all acceptable gambles.

TABLE 2

Preference Gambles through Which Payoff Functions Are Revealed

Outcome	$(\mathbf{u}_T - \mathbf{u}_B)\mathbf{1}_T$	$(\mathbf{u}_B - \mathbf{u}_T)\mathbf{1}_B$	$(\mathbf{u}_L - \mathbf{u}_R)\mathbf{1}_{aL}$	$(\mathbf{u}_L - \mathbf{u}_R)\mathbf{1}_{bL}$	$(\mathbf{u}_R - \mathbf{u}_L)\mathbf{1}_{aR}$	$(\mathbf{u}_R - \mathbf{u}_L)\mathbf{1}_{bR}$
aTL	1	0	1	0	0	0
aBL	0	−1	1	0	0	0
aTR	−1	0	0	0	−1	0
aBR	0	1	0	0	−1	0
bTL	1	0	0	−1	0	0
bBL	0	−1	0	−1	0	0
bTR	−1	0	0	0	0	1
bBR	0	1	0	0	0	1

TABLE 3

Belief Gambles through Which Prior Probabilities Are Revealed

Outcome	$(\mathbf{1}_a - p_a^*)\mathbf{1}_T$	$(\mathbf{1}_a - p_a^*)\mathbf{1}_B$	$(\mathbf{1}_b - p_b^*)\mathbf{1}_T$	$(\mathbf{1}_b - p_b^*)\mathbf{1}_B$
aTL	0.4	0	−0.4	0
aBL	0	0.4	0	−0.4
aTR	0.4	0	−0.4	0
aBR	0	0.4	0	−0.4
bTL	−0.6	0	0.6	0
bBL	0	−0.6	0	0.6
bTR	−0.6	0	0.6	0
bBR	0	−0.6	0	0.6

An acceptable gamble is any gamble which dominates a gamble of the form $A\alpha + B\beta$, where α and β are nonnegative vectors.

3. Joint Coherence

The rationality of the players' behavior in the game will now be analyzed from the perspective of an outside observer ("she") who is completely naive concerning the reputations of the players, their stakes in the game, or the probabilities of states of nature. Suppose that the observer chooses to enforce some of the gambles the players have offered to accept at the conclusion of their preplay communication. Then the criterion for a "rational" outcome of the game is quite simple: the observer should not succeed in extracting money from the players without putting any money at risk. In other words, she should not be able to find an acceptable gamble which yields a nonpositive aggregate payoff to the players under every outcome of the game and which yields a strictly negative aggregate payoff under the outcome actually observed. Such a gamble constitutes an *arbitrage opportunity* against the observed outcome. To formalize this criterion, let (s^*, t^*) denote the observed outcome and let $\mathbf{1}_{s^*,t^*}$ denote the corresponding indicator vector. Then we assume:

A6 (No arbitrage opportunities). $-\mathbf{1}_{s^*,t^*} \notin \mathcal{A}$.

DEFINITION. The outcome (s^*, t^*) is *jointly coherent* if A6 holds given A1–A5—*i.e.*, if there do not exist (α, β) such that $[A\alpha + B\beta](s, t) \leq 0$ for all $(s, t) \in S \times T$ and $[A\alpha + B\beta](s^*, t^*) < 0$.

Notice that A6 is a joint restriction on the set $\mathcal{A}$ and the outcome (s^*, t^*): the strategy chosen by the players given their information must "cohere" with the beliefs they have previously revealed about the game's structure. In the case of a nonstrategic decision problem ($N = 1$), this is merely de Finetti's standard of individual rationality.[7] In the strategic case ($N > 1$), it captures the intuitive idea of *mutually expected* rationality, namely that the players should not behave incoherently as individuals, nor bet on each other to behave incoherently, nor bet on each other to bet on each other to behave incoherently, and so on (Nau and McCardle 1990). Necessary and sufficient conditions for joint coherence are summarized in:

THEOREM 1. *The outcome (s^*, t^*) is jointly coherent if and only if there exists a probability distribution π on $S \times T$ which*:

(i) *assigns positive probability to (s^*, t^*)*,

(ii) *is a supporting probability distribution for every player*; *and*

[7] Actually, in the case where $N = 1$, A6 is a slight strengthening of de Finetti's standard: it requires the avoidance of arbitrage opportunities *ex post*, not merely *ex ante*. The *ex post* version is appropriate in the presence of information and control over events. Here, the players jointly know t^* at the time they jointly choose s^*, so we can say "they should have known better or else acted otherwise" if arbitrage is possible against (s^*, t^*).

(iii) *has the property that, if it is interpreted as the joint distribution of types and "recommended" strategies for the players,*[8] *then each player's recommended strategy maximizes his expected payoff given his information and given that the other players follow their own recommendations—i.e.*:

$$P_\pi(\mathbf{u}_{nj} \mid \mathbf{1}_{njr}) \geq P_\pi(\mathbf{u}_{nk} \mid \mathbf{1}_{njr}) \qquad \textit{or else} \qquad P_\pi(\mathbf{1}_{njr}) = 0 \qquad \forall j, k \in S_n, \qquad r \in T_n.$$

PROOF. By linear duality (*e.g.*, Gale 1960, Theorem 3.10), either there exist $\alpha \geq 0$, $\beta \geq 0$ such that $A\alpha + B\beta \leq -\mathbf{1}_{s^*,t^*}$, or else there exists $\pi \geq 0$ such that: (i) $\pi^T \mathbf{1}_{s^*,t^*} > 0$, (ii) $\pi^T B \geq 0$, and (iii) $\pi^T A \geq 0$. Then the "either" condition is the existence of (α, β) constituting an arbitrage opportunity against (s^*, t^*); and the "or else" condition is the existence of a π satisfying parts (i), (ii), and (iii), respectively, of the theorem. □

The distribution π can be factored as $\pi(s, t) = \rho(t)\mu(s \mid t)$, where ρ is what Harsanyi (1967) calls a *basic probability distribution of the game* (a "common prior" distribution); and μ is a *correlated Bayesian equilibrium distribution*. The latter can be viewed as a coordination mechanism used by a mediator for generating self-enforcing randomized strategy recommendations (possibly correlated between players) conditioned on states of nature. Such a mediator may be omniscient—*i.e.*, know the true state of nature.

The theorem can be restated more informally as follows: joint coherence requires the players to act *as if* they held *some* common prior distribution over states of nature, consistent with their revealed beliefs, and employed a coordination mechanism (possibly requiring their true information as input) from whose recommendations they would not have incentive to deviate unilaterally given their information. Joint coherence is thus seen to be dual to a correlated analog of Harsanyi's Bayesian equilibrium concept[9] for games of incomplete information, just as it is dual to a correlated analog of Nash equilibrium (namely, correlated equilibrium) in the complete-information case. Once again, joint coherence does not presume that distributions (over states and/or strategies) are uniquely determined. Rather, the game must be played only *as if some* such distributions existed.

As an illustration of Theorem 1, the game introduced in §2 has exactly two outcomes which are jointly coherent, namely aTL and bTR. The unique supporting common prior distribution is $\rho(a) = 0.6$, $\rho(b) = 0.4$, and the unique supporting strategy-recommendation mechanism is $\mu(TL \mid a) = \mu(TR \mid b) = 1$. This is also a Bayesian equilibrium in the sense of Harsanyi because the recommended strategies of the players are (trivially) independent. This occurs because the belief gambles constituting the B matrix indicate that the players believe independence to hold. If the players had *not* believed independence to hold, they could have indicated this by accepting only the two unconditional gambles $\mathbf{1}_a - p_a^*$ and $\mathbf{1}_b - p_b^*$ rather than the four conditional gambles $(\mathbf{1}_a - p_a^*)\mathbf{1}_T$, $(\mathbf{1}_a - p_a^*)\mathbf{1}_B$, $(\mathbf{1}_b - p_b^*)\mathbf{1}_T$ and $(\mathbf{1}_b - p_b^*)\mathbf{1}_B$. In this case, the outcomes aBL and bBR would also have been jointly coherent, supported (for example) by the correlated strategy-recommendation mechanism $\mu(TL \mid a) = \frac{1}{3}$, $\mu(BL \mid a) = \frac{2}{3}$, and $\mu(BR \mid b) = 1$. However, the latter mechanism could be implemented only by an omniscient mediator, since it gives player 2 an incentive to lie in order to induce player 1 to play T in state b.

This example illustrates an important difference between joint coherence and alternative concepts such as rationalizability (Bernheim 1984, Pearce 1984) or Nash equilibrium and its refinements: assumptions about independence are embedded not in an objective standard of rational play, but rather in the subjectively revealed rules of the game. The

[8] If π is interpreted as a distribution for jointly determining types and recommended strategies, then $P_\pi(\mathbf{1}_{njr})$ is the joint probability that player n will be of type r and receive recommendation j, and $P_\pi(\mathbf{u}_{nk} \mid \mathbf{1}_{njr})$ is player n's conditional expected payoff for playing strategy k given that his type is r and he has received recommendation j, given that other players follow their own recommendations.

[9] Harsanyi's concept requires the strategy-generating mechanism to maintain statistical independence between the players given their information—i.e., it must have the factorization $\mu(s \mid t) = \mu_1(s_1 \mid t_1) \times \cdots \times \mu_N(s_N \mid t_N)$.

players are expected to act independently only to the extent they reveal that they *believe* their actions will be independent. We do not rule out *a priori* the possibility that they might correlate their actions through direct communication or through employment of a possibly-omniscient mediator. (Omniscient mediation may be possible if, for example, there exist institutional mechanisms for punishing untruthful reporting of types.) In the example just discussed, there would be no advantage to either player in employing a mediator, but the payoff structure could be modified so as to create such an advantage. For example, if the payoff to player 2 in outcome aBL were raised from 4 to 5, a stochastically dominant unconditional payoff distribution could be achieved for player 2 by employing an omniscient mediator.

The criterion of joint coherence can be applied even in the absence of uniquely determined probabilities or payoff functions. For example, the common prior may be indeterminate: with $p_a^* = 0.5$ and $p_b^* = 0.4$ in the example of §2, the outcomes aBL and bBR would also be jointly coherent, supported by an alternative prior distribution $\rho(a) = \rho(b) = 0.5$ and coordination mechanism $\mu(BL|a) = \mu(BR|b) = 1$.

4. Games with Observable Communication

If the players desire to coordinate their actions but omniscient mediation is believed to be impossible, this must somehow be written into the rules of the game before it can be asserted that outcomes supportable only by omniscient guidance are irrational. One way to do this is to suppose that the players employ a mechanical communication device with commonly known properties. This idea has been used by numerous authors (Myerson 1985, 1986; Farrell 1985; Forges 1986) to extend Harsanyi's model of a Bayesian game, allowing strategies to be correlated in ways which are *incentive compatible—i.e.*, which encourage truthful reporting of types and obedience to recommended strategies. When the Nash solution concept is applied to the extended game, it yields the concept of a *communication equilibrium*: a correlated equilibrium that could in principle be implemented by a nonomniscient mediator who relies on the players to report their private information. Thus, by adding a communication mechanism to the game, the set of Bayesian equilibrium outcomes can be enlarged to include outcomes supported by correlated strategies which would otherwise be ruled out by the Nash assumption of statistical independence.

Here, the concept of communication equilibrium will be approached from the opposite direction: by introducing a communication mechanism into the rules of the game, the set of jointly coherent outcomes can be restricted to those which are achievable without omniscient guidance. The incorporation of a formal communication mechanism does not obviate the need for preplay communication, but merely removes it to a higher level. Indeed, more elaborate unobserved communication is needed in this setting to achieve common knowledge of the rules, because the extended game has a much larger strategy set than the original game. To establish the duality between joint coherence and communication equilibrium in a game with observable communication, an objective description of a communication game will first be given, and it will then be shown how its structure can be subjectively revealed through the acceptance of gambles.

The objective description of a game with observable communication begins with a strategy set $S = S_1 \times \cdots \times S_N$, a set of states of nature $T = T_1 \times \cdots \times T_N$, and payoff functions $(u_1, \ldots, u_N)$. Add to this a deterministic communication mechanism $m = (m_1, \ldots, m_N)$ which receives inputs $i = (i_1, \ldots, i_N)$ from the players and returns outputs $o = (o_1, \ldots, o_N)$ to them, where $o_n = m_n(i)$. (The output to player n may depend on the inputs of all players.) The sets of inputs $\{i_n\}$ and outputs $\{o_n\}$ available to player n will be denoted I_n and O_n, respectively, with $I = I_1 \times \cdots \times I_N$ and $O = O_1 \times \cdots \times O_N$. It is usually assumed that each player has at least one input for each of his possible states of information and at least one output for each of his available strategies.

In the presence of such a mechanism, the players' choices of strategies in the original game are replaced by choices of *decisions*, where a decision d_n for player n consists of a function $f_n : T_n \mapsto I_n$ mapping information states into inputs, and a function $g_n : T_n \times O_n \mapsto S_n$ mapping information states and outputs into strategies. Henceforth, let $d_n = (f_n, g_n)$ denote a decision of player n; let $d = (d_1, \ldots, d_N)$ denote a joint selection of decisions by all players; let d_{-n} denote a joint selection of decisions by all players other than n; and let D and D_{-n} denote the sets of all $\{d\}$ and $\{d_{-n}\}$, respectively. Let $f(t) = (f_1(t_1), \ldots, f_N(t_N))$, and let $s^*(d, m, t)$ denote the strategy in the original game which results from the players' choice of decision d and the selection of mechanism m when the state of nature is t: that is, $s_n^*(d, m, t) \equiv g_n(t_n, m_n(f(t)))$. The payoff to player n as a function of d, t, and m is then $u_n(s^*(d, m, t), t)$.

To allow randomization of strategies, either independently or with correlation, the choice of the communication mechanism is allowed to be objectively randomized. Let M denote the set of all deterministic mechanisms, *i.e.*, the set of all mappings $\{m : I \mapsto O\}$. Then let a particular mechanism m be chosen from M by objective randomization using a known distribution μ, independent of s and t, so that the identity of m is unknown to the players at the time they give their inputs or receive their outputs. Henceforth, the distribution μ will itself be referred to as the (random) mechanism employed by the players.

The objective of introducing observable communication is to enable the players to implement a self-enforcing, correlated strategy which does not require an omniscient mediator, presumably in order to achieve a mutually desired distribution of payoffs which would otherwise be unreachable in a noncooperative setting. It may therefore be assumed that the players wish to avoid vagueness, and that their preplay communication accordingly leads to (among other things) the selection of a unique mechanism μ^* and a unique joint decision d^* to be employed with it. (Of course, the joint decision cannot be enforced; it is merely proposed as a focal point. The question is whether it is rational even to assert the intention of adhering to decision d^* with respect to mechanism μ^*.) Without loss of generality, the joint decision may be assumed to be deterministic, since a randomized decision with respect to one mechanism would be equivalent to a deterministic decision with respect to some other mechanism.[10] For all n and all $k \in D_n$, let $d^*_{-n}k = (d_1^*, \ldots, d_{n-1}^*, k, d_{n+1}^*, \ldots, d_N^*)$ and let $\mathbf{u}_{nk}^*$ denote the vector whose elements, indexed by states of nature, are the *expected* payoffs to player n when $d^*_{-n}k$ is implemented. That is, $\mathbf{u}_{nk}^*$ is the expected payoff to player n when he implements decision k while his opponents implement d^*_{-n}, where expectation is with respect to the distribution μ^*. Formally:

$$u_{nk}^*(t) \equiv \sum_{m \in M} \mu^*(m) u_n(s^*(d^*_{-n}k, m, t), t).$$

The *subjective* description of the communication game can now be given. Let $\mathbf{1}_{nr}$ denote the indicator function defined on T alone (rather than $S \times T$) for the event that player n is of type r. Then let information about the players' expected payoffs be revealed through the acceptance of preference gambles defined on T in the following way:

A4′ (Acceptance of preference gambles). For every n, every $k \in D_n$, every $r \in T_n$, and $j = d_n^* : (\mathbf{u}_{nj}^* - \mathbf{u}_{nk}^*)\mathbf{1}_{nr} \in \mathcal{A}$.

In other words, given that his type is r, player n prefers the *expected* payoffs he obtains by adhering to decision $j = d_n^*$ to those he would obtain by defecting to decision k, given that his opponents are adhering to d^*_{-n}. (The condition that his opponents adhere to

[10] Indeed, by the same reasoning, we could assume that μ^* is a "canonical" mechanism for which I_n is a copy of T_n and O_n is a copy of S_n for all n, and that d^* is the unique decision in which each player is truthful and obedient—i.e., $f_n^*(t_n) = t_n$ and $g_n^*(t_n, o_n) = o_n$ for all n, t_n and o_n. This ability to restrict attention to canonical mechanisms and truthful-and-obedient behavior is known as the "revelation principle" (Myerson 1985, Forges 1986).

d^*_{-n} means that the payoffs in the gamble depend on what the other players *would have done* as a function of their information if they had adhered, not on the decisions they actually implemented if they defected.) If A4′ holds, it is *as if* the outcome set of the game is $S \times T$, the payoff functions are $(u_1, \ldots, u_N)$, and decision d^* is employed in conjunction with mechanism μ^*, and this is common knowledge. Notice, however, that the details of μ^* and d^* (and, indeed, the details of the sets I and O) have been integrated out of the description of the game insofar as the observer is concerned: the payoffs between the players and the observer depend only the state t, not on d, s, or m. This is a consequence of the revelation principle, which establishes that many different combinations of mechanisms and decisions are equivalent, and the fact that we have integrated over m to operationalize the players' agreement on the distribution μ^*. The situation is analogous to Aumann's (1987) formulation of a game of complete information, in which the pegging of strategies to exogenous states is used as a means to achieve intra-play communication.

Let A henceforth denote the matrix whose column in the (n, k, r) position is the preference-gamble payoff vector $(\mathbf{u}^*_{nj} - \mathbf{u}^*_{nk})\mathbf{1}_{nr}$, where $j = d^*_n$, and whose rows are indexed (only) by t. α will henceforth denote the corresponding vector of nonnegative gamble coefficients chosen by an observer. Also, since the players are pegging their strategies to their types through a known mechanism and decision, it will be assumed henceforth that the events $\{\mathbf{E}_{nh}\}$ and $\{\mathbf{F}_{nh}\}$ to which belief gambles refer are measurable with respect to T (rather than $S \times T$). The matrix B, whose columns are the payoff vectors for the belief gambles, will now also have rows indexed by t. Assumption A5 (acceptable belief gambles) is otherwise retained. Under these new definitions, the set $\mathcal{A}$ of acceptable gambles is once again the set of all gambles that equal or dominate a gamble of the form $A\alpha + B\beta$, where α and β are nonnegative vectors.

As before, the standard of rationality applied to the players is that an observer should not succeed in winning money from them without having put money at risk. In this case, the outcome of the extended game is t^*, the observed vector of types, insofar as the observer is concerned. We therefore modify A6 to require that there should be no acceptable gamble which yields a strictly negative aggregate payoff in state t^* and nonnegative aggregate payoffs in all other states:

A6′ (No arbitrage opportunities). $-\mathbf{1}_{t^*} \notin \mathcal{A}$.

DEFINITION. The outcome of the communication game is *jointly coherent* if A6′ is satisfied given A1-A2-A3-A4′-A5—*i.e.*, if there do not exist (α, β) such that $[A\alpha + B\beta](t) \leq 0$ for all $t \in T$ and $[A\alpha + B\beta](t^*) < 0$.

THEOREM 2. *The outcome of the communication game is jointly coherent if and only if there exists a distribution ρ on T which*:

(i) *assigns positive probability to t^**;

(ii) *is consistent with the asserted beliefs p^**; *and*

(iii) *has the property that each player's expected payoff is maximized by adhering to decision d^* with respect to the mechanism μ^*, given that all the other players do likewise.*

PROOF. By linear duality, either there exist $\alpha \geq 0$, $\beta \geq 0$ such that $A\alpha + B\beta \leq -\mathbf{1}_{t^*}$, or else there exists a distribution ρ such that $\rho^T A \geq 0$, $\rho^T B \geq 0$, and $\rho^T \mathbf{1}_{t^*} > 0$. The system of inequalities $\rho^T A \geq 0$ can be written out in full as:

$$\sum_{t \in T} \sum_{m \in M} \rho(t)\mu^*(m)[u_n(s^*(d^*, m, t), t) - u_n(s^*(d^*_{-n}k, m, t), t)] \geq 0 \qquad (2)$$

for all n and all $k \in D_n$, which expresses the condition that the expected payoff to player n is greater for adhering to decision d^*_n than for defecting to decision k, given that the other players adhere. The system $\rho^T B \geq 0$ can be written as:

$$\sum_{t \in T} \rho(t)[E_{nh}(t) - p^*_{nh}]F_{nh}(t) \geq 0 \qquad \forall n \text{ and } h \in H_n.$$

Letting $P_\rho(\cdot)$ denote the probability measure on subsets of T induced by ρ, this means that either $P_\rho(\mathbf{F}_{nh}) = 0$ or else $P_\rho(\mathbf{E}_{nh} | \mathbf{F}_{nh}) \geq p^*_{nh}$ for every n and h. That is, the distribution ρ is consistent with the announced lower probabilities p^*. Finally, the inequality $\rho^T\mathbf{1}_{t^*} > 0$ is equivalent to $\rho(t^*) > 0$. □

Thus, joint coherence requires the existence of a distribution on states of nature under which μ^* in conjunction with d^* constitutes a communication equilibrium as defined by Forges (1986) and Myerson (1985, 1986). Moreover, this distribution must agree with the announced beliefs and must assign positive probability to the state of nature that is actually observed. Theorem 2 thus supports the concept of communication equilibrium in games with observable communication, absent the assumption that the basic distribution of the game is fully revealed. In particular, equation (2) expresses the condition that the mechanism μ^* in conjunction with the decision d^* must be incentive compatible with respect to the distribution ρ. (This is equivalent to equation (5.4) in Myerson 1985.) Forges (1986) has given a characterization of the set of expected payoffs achievable by all mechanisms and decisions yielding communication equilibria with respect to a given state-distribution ρ: this set of expected payoffs is a convex polyhedron, being defined by a system of linear inequalities. Theorem 2 inversely characterizes the set of all state-distributions $\{\rho\}$ which support a given mechanism and decision and are also consistent with whatever beliefs have been announced: this set of distributions is also a convex polyhedron, being defined by the systems of inequalities given above, which are bilinear in ρ and μ^*. For fixed ρ, the inequalities always have a canonical (*i.e.*, truthful-and-obedient) solution in μ^*, since there is always at least one Bayesian equilibrium: a "non-communicative" communication equilibrium. However, for fixed μ^* and d^*, there need not be a solution in ρ. This could occur, for example, if d^* prescribed strategy choices which were individually incoherent.

5. Discussion

It has been shown that the solution concepts of (correlated) Bayesian equilibrium and communication equilibrium, and the related concepts of common prior distributions and incentive-compatible mechanisms, can be derived from a subjective model of a noncooperative game. In this model, the players elicit each other's probabilities and payoffs by the operational method of de Finetti, and strategic rationality is defined as the avoidance of arbitrage opportunities against the group. Closely related arbitrage arguments also characterize competitive equilibria in securities markets and exchange economies (Nau and McCardle 1991). These results suggest that there is a deeper unity among subjective probability theory, noncooperative game theory, and competitive market theory than is commonly appreciated—a unity that does not depend on explicit or highly determinate cognitive models. They also suggest that, in the application of game-theoretic reasoning to economic decision-making, the preplay communication process is an essential object of study. Without a model of preplay communication, there is no sound basis for asserting that common knowledge exists, nor for predicting or prescribing how the players should choose among multiple strategies that may meet the requirement of mutually expected rationality.

Throughout this paper, the simplifying assumption has been made that players' marginal utilities for monetary wealth are constant across outcomes of the game, in which case their "true" probabilities and relative utilities are in principle directly revealed by their acceptance of monetary gambles. Under the more realistic assumption of risk-averse (concave) utility for money, state-dependence of wealth will imply state-dependence of marginal utilities, leading to distortions of gambling-based measurements of probability and utility. The apparent probabilities revealed by a player's acceptance of belief gambles must in this case be interpreted as renormalized products of his "true" probabilities and his marginal utilities, as pointed out by Kadane and Winkler (1988). This phenomenon is in fact central to the argument that communication-through-gambling will produce convergence to an apparent common prior distribution (Nau 1990). Its implications for coherent behavior in games are analyzed in Nau (1991b), where it is shown that the revealed utilities of risk-averse players are distorted in a fashion reciprocal to that of their probabilities: the players' apparent utility differences, as revealed by their acceptance of preference gambles, must be interpreted as their true utility differences divided by their marginal utilities. Fortunately, these reciprocal distortions cancel out when expected utilities are computed, so that it remains valid to use the revealed probabilities and utilities in determining rational outcomes of the game. The rational outcomes are those which occur with positive probability in an *objective* correlated equilibrium of the revealed game (in which it is common knowledge that the players hold the same prior distribution on states), and this corresponds to a *subjective* correlated equilibrium of the underlying "true" game (in which it is common knowledge that the players hold different priors—*cf.* Aumann 1974). This establishes a relativity between objective and subjective correlated equilibria as expressions of strategic rationality: the former are what can be observed, the latter are what may be imagined as the "truth."[11]

[11] The author is grateful for comments on earlier drafts of this paper by Francoise Forges, Doug Foster, Jim Friedman, Dan Graham, Kevin McCardle, Hervé Moulin, David Schmeidler, Peter Wakker, Bob Winkler, the editors, and several anonymous referees. The opinions expressed herein and any remaining errors are the sole responsibility of the author. This research was supported by the Business Associates Fund at the Fuqua School of Business.

References

ARMBRUSTER, W. AND W. BÖGE, "Bayesian Game Theory," In O. Moeschlin and D. Pallaschke (Eds.), *Game Theory and Related Topics*. North-Holland, Amsterdam, 1979.

AUMANN, R., "Utility Theory Without the Completeness Axiom," *Econometrica*, 30 (1962), 445–462.

———, "Subjectivity and Correlation in Randomized Games," *J. Math. Economics*, 1 (1974), 67–96.

———, "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55 (1987), 1–18.

BERNHEIM, D. B., "Rationalizable Strategic Behavior," *Econometrica*, 52 (1984), 1007–1028.

BÖGE, W. AND T. EISELE, "On Solutions of Bayesian Games," *Internat. J. Game Theory*, 8 (1979), 193–215.

DE FINETTI, B., "Foresight: Its Logical Laws, Its Subjective Sources," (1937). Translation reprinted in: H. E. Kyburg and H. E. Smokler (Eds.), *Studies in Subjective Probability*, Robert Krieger, New York, 1980, 53–118.

———, *Theory of Probability*. Vol. 1, Wiley, New York, 1974.

ELLSBERG, D., "Risk, Ambiguity, and the Savage Axioms," *Quart. J. Economics*, 75 (1961), 643–669.

FARRELL, J., "Communication Equilibrium in Games," GTE Laboratories Economics Working Paper 85-3, 1985.

FISHBURN, P. C., "Nontransitive Measurable Utility," *J. Math. Psychology*, 26 (1982), 31–67.

FORGES, F., "An Approach to Communication Equilibria," *Econometrica*, 54 (1986), 1375–1385.

GALE, D., *Theory of Linear Economic Models*, McGraw-Hill, New York, 1960.

HARSANYI, J., "Games with Incomplete Information Played by "Bayesian" Players. I–III," *Management Sci.*, 14 (1967), 3, 5, 7 159–182, 320–334, 486–502.

———, "Subjective Probability and the Theory of Games: Comments on Kadane and Larkey's Paper, and Rejoinder to Reply," *Management Sci.*, 28, 2 (1982), 120–124, 124–125.

KADANE, J. B. AND P. D. LARKEY, "Subjective Probability and the Theory of Games, and Reply," *Management Sci.*, 28, 2 (1982), 113–120, 124.

——— AND ———, "The Confusion of Is and Ought in Game Theoretic Contexts," *Management Sci.*, 29, 12 (1983), 1365–1379.

——— AND R. WINKLER, "Separating Probability Elicitation from Utilities," *J. Amer. Statist. Assoc.*, 83, (1988), 402 357–363.

KAHAN, J., "On Choosing Between Rev. Bayes and Prof. von Neumann," *Management Sci.*, 29, 11 (1983), 1334–1336.

LEAMER, E., "Bid-Ask Spreads for Subjective Probabilities," In P. Goel and A. Zellner (Eds.), *Bayesian Inference and Decision Techniques*, Elsevier, New York, 1986, 217–232.

MACHINA, M., "Expected Utility Without the Independence Axiom," *Econometrica*, 50 (1982), 227–323.

MERTENS, J. F. AND S. ZAMIR, "Formulation of Bayesian Analysis for Games with Incomplete Information," *Internat. J. Game Theory*, 14 (1985), 1–29.

MYERSON, R., "Bayesian Equilibrium and Incentive Compatibility," In L. Hurwicz, D. Schmeidler, and H. Sonnenschein (Eds.), *Social Goals and Social Organization*, Cambridge University Press, Cambridge, 1985, 229–259.

———, "Multistage Games with Communication," *Econometrica,* 54 (1986), 323–358.

NAU, R., "The Incoherence of Agreeing to Disagree," Fuqua School of Business Working Paper 9006, 1990.

———, "Coherent Decision Analysis with Inseparable Probabilities and Utilities," Fuqua School of Business Working Paper 9101, 1991a.

———, "The Relativity of Subjective and Objective Correlated Equilibria," Fuqua School of Business Working Paper 9102, 1991b.

———, "Indeterminate Probabilities on Finite Sets," *Ann. Statist.*, (1992) (forthcoming).

——— AND K. MCCARDLE, "Coherent Behavior in Noncooperative Games," *J. Economic Theory*, 50, 2 (1990), 424–444.

——— AND ———, "Arbitrage, Rationality, and Equilibrium," *Theory and Decision*, 31 (1991), 199–240.

PEARCE, D., "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 54 (1984), 1029–1050.

ROTH, A. AND F. SCHOUMAKER, "Subjective Probability and the Theory of Games: Some Further Comments," *Management Sci*., 29, 11 (1983), 1337–1340.

ROTHKOPF, M., "Modeling Semirational Competitive Behavior," *Management Sci*., 29, 11 (1983), 1341–1346.

SCHMEIDLER, D., "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57 (1989), 571–587.

SEIDENFELD, T., "Decision Theory Without "Independence" or Without "Ordering"," *Economics and Philosophy*, 4 (1988), 267–315.

SHUBIK, M., "Comment on the "The Confusion of Is and Ought in Game-Theoretic Contexts"," *Management Sci*., 29, 12 (1983), 1380–1383.

SMITH, C., "Consistency in Statistical Inference and Decision," *J. Roy. Statist. Soc. Ser. B*, 23 (1961), 1–25.

TAN, T. AND S. WERLANG, "The Bayesian Foundations of Solution Concepts of Games," *J. Economic Theory*, 45 (1988), 370–391.

VON NEUMANN, J. AND O. MORGENSTERN, *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, NJ, 1944.